



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Distributed Multisensor Fusion System Specification and Evaluation Issues

Jonathan A. Legg

**Intelligence, Surveillance and Reconnaissance Division
Defence Science and Technology Organisation**

DSTO-TN-0663

ABSTRACT

This report discusses the tracking and sensor fusion issues that should be addressed during the specification or evaluation of a distributed, multisensor surveillance system for the provision of a common situational awareness picture or remote weapons control data. These issues include low level sensor data, sensor registration, target state estimation, sensor networking, sensor control, hardware and software, and performance specification and assessment.

APPROVED FOR PUBLIC RELEASE

Published by

*Defence Science and Technology Organisation
PO Box 1500
Edinburgh, South Australia 5111, Australia*

*Telephone: (08) 8259 5555
Facsimile: (08) 8259 6567*

*© Commonwealth of Australia 2005
AR No. 012-918
October, 2005*

APPROVED FOR PUBLIC RELEASE

Distributed Multisensor Fusion System Specification and Evaluation Issues

EXECUTIVE SUMMARY

The increased use of multiple data and information sources has led to the development of distributed surveillance systems, in which multiple sensors, platforms and information sources exchange information over one or more networks. The information is used to form a common picture of the volume or region of interest at the various nodes throughout the system for the purposes of situational awareness and weapon control. The nodes of a distributed system must be able to handle plot and/or track data of various content and format from a diverse range of sensors.

By virtue of the increased system complexity, additional low level requirements, specifications and test routines are required to support the design, development and acceptance into service of distributed surveillance systems compared to single platform systems. This technical note addresses a number of the issues that must be considered when specifying and evaluating a distributed surveillance system. The specific areas addressed include sensor data processing, sensor registration, track data processing, sensor networking, sensor control, hardware and software issues, and performance specification and assessment, with the specific issues within each area discussed in detail. The relative priority or importance of each issue is dependent on the specific application and system.

Contents

Glossary	ix	
1	Introduction	1
2	Multisensor Fusion Architectures	1
3	Distributed Multisensor Fusion Issues	5
4	Conclusions	19
5	Acknowledgements	19
References	19	

Appendices

A	Summary of Issues	21
----------	--------------------------	-----------

Figures

1	Platforms sharing a common picture of the environment.	2
2	Example distributed sensor fusion architecture.	3
3	Sensor architecture.	4
4	Tracker architecture.	4

Tables

A1	Summary of distributed multisensor fusion issues.	21
----	---	----

Glossary

AMR	Associated Measurement Report
C2	Command and Control
CEC	Cooperative Engagement Capability
GPS	Global Positioning System
JDL	Joint Directors of Laboratories
MOP	Measure of Performance
R2	Reporting Responsibility
TADIL	Tactical Digital Information Link
TQ	Track Quality

DSTO-TN-0663

1 Introduction

Modern civilian and military surveillance systems are moving from single platform systems to distributed systems that use data and information from multiple platforms and sources. Such systems require a network or networks to distribute the information, and more complex algorithms for processing the information available at each node in the network. The information processing algorithms must be able to register, correlate, associate and fuse information from a variety of diverse sources. This extra design complexity necessitates additional low level requirements, specifications and test routines to support the system design, development and acceptance into service.

A typical requirement of a distributed surveillance system is to maintain a common and accurate representation of the position, motion and identity of all the vehicles within the surveillance region. This information is typically presented to a user for situational awareness via a Command and Control (C2) system, or for other purposes, including the control of weapons. Such systems allow commanders on various participating platforms to share a common view of the environment and coordinate their actions. Examples of quite different data distribution systems are the Tactical Digital Information Link (TADIL) and the United States' Cooperative Engagement Capability (CEC).

Due to the intricate melding of different types of sensors, communications systems, signal processing algorithms, data fusion processes, manufacturers and end user applications, distributed surveillance systems are inherently complex and have a high level of risk associated with their procurement. There are many aspects to their performance, and the behaviour of each item of equipment needs to be specified or understood with regard to its effects on other systems.

The objective of this document is to identify and describe those aspects of a distributed, multisensor tracking system that should be considered when specifying or evaluating such a system, or the components within such a system. It does not discuss the relative merits of the possible network architectures, the higher levels of data fusion, such as situation awareness and impact assessment [Hall & Llinas 2001], or the use of the resulting data. It is assumed that one or more sensors are associated with a platform that processes the sensor data and shares that data, or its products, with a number of other platforms. The remainder of this document is structured as follows. Section 2 describes the multisensor fusion architectures addressed in this report, a discussion of the major issues is presented in Section 3, and Section 4 is the conclusion.

2 Multisensor Fusion Architectures

Figure 1 shows a distributed, multisensor data fusion scenario. Several platforms, each with sensors and communications systems, share data and form a common picture of the environment. There may be communications with other, external participants. This scenario represents a number of communicating surface combatants or airborne early warning and control aircraft, or a combination of these, for example. The sensor fusion issues addressed by this report are at the 'object assessment' level, or Level 1, in the revised US Joint Directors of Laboratories (JDL) data fusion model [Steinberg & Bowman 2001].

The product of interest is a common, accurate track picture available to the platforms' C2 systems.

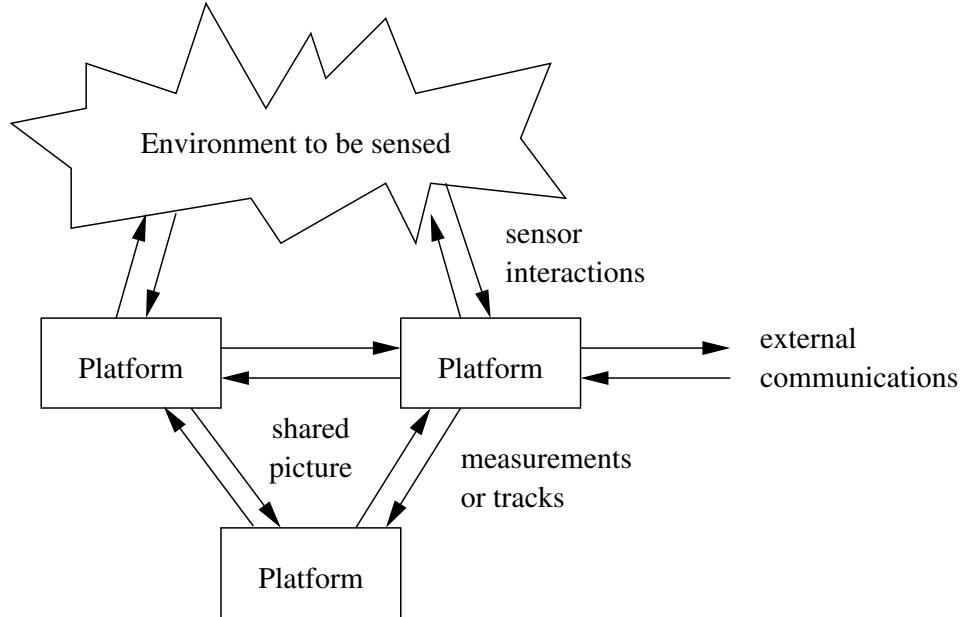


Figure 1: Platforms sharing a common picture of the environment.

The tracking, data fusion and communications elements may be organised in many different ways, such as the architecture shown in Figure 2. This shows the major elements of a platform, which acts as a data gathering and processing node. A radar, for example, provides measurements to a local tracker, which provides low level control over the radar, and utilises global tracks, when available, for plot-to-track association. This approach helps maintain consistency between the local and global track databases¹. Local tracks, or measurements associated with local tracks, are passed to a global tracking/sensor fusion system and other platforms' sensor fusion systems via a high bandwidth data link. The output from the local and global trackers are passed into the C2 system, which provides a higher level of communications with other platforms via a TADIL. It also provides communications with other entities and high level control over the radar (such as cuing). Details of the sensor and tracker blocks shown in Figure 2 are given in Figures 3 and 4, respectively. The sensor subsystems may have inbuilt trackers, or simply return measurements. The degree of external control depends upon the nature of the sensor. Trackers manage track databases, forming new tracks and dropping old ones, and associate and combine incoming measurements with existing tracks to provide target state estimates and predictions.

In practice, a high level of complexity exists within a system such as this, resulting from the interaction of the many different subsystems, each of which are complex in their own right. For example, the 'C2 system' in Figure 2 needs to resolve many issues, including the consistency of the track pictures provided by the many information sources.

¹Conversation with Dr Andrew Shaw, 2003.

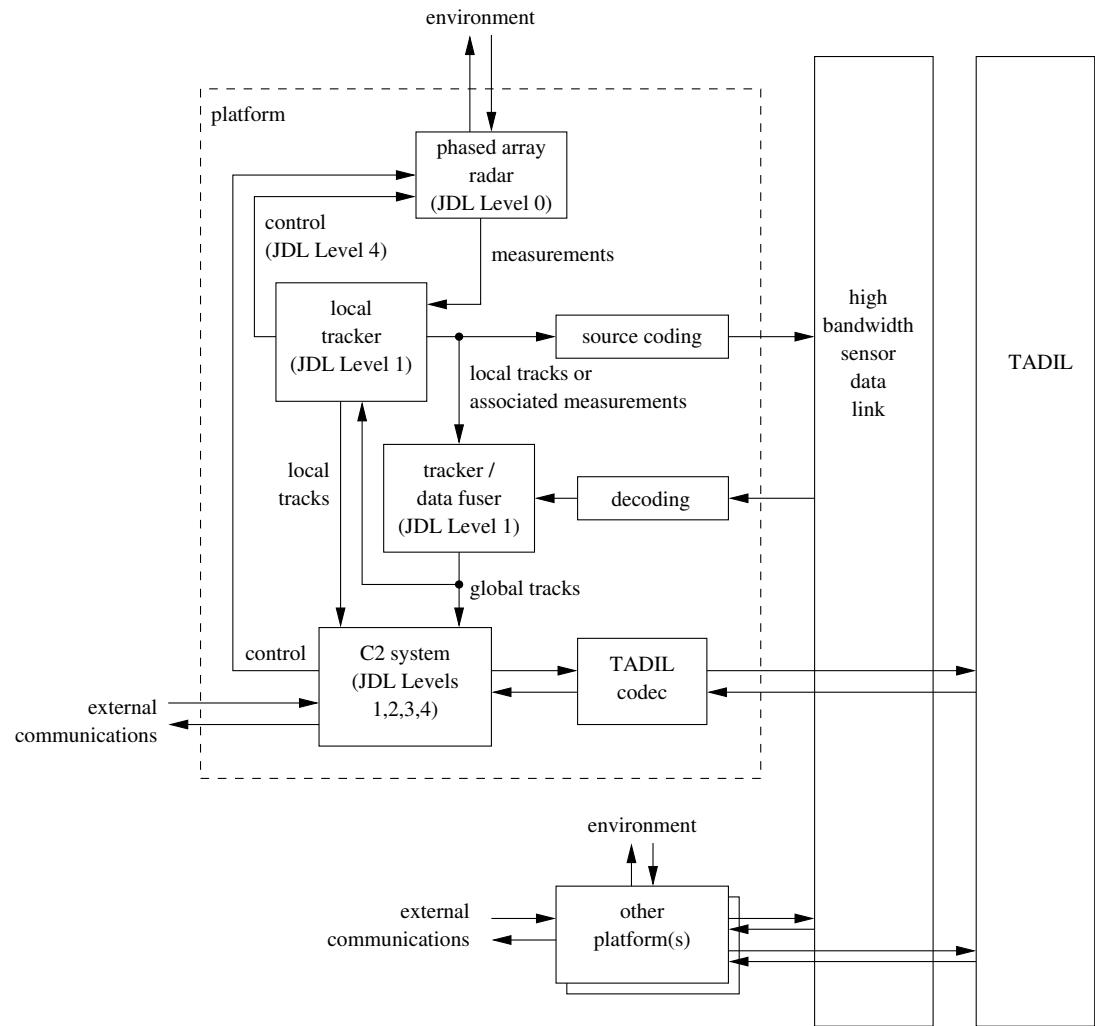


Figure 2: Example distributed sensor fusion architecture.

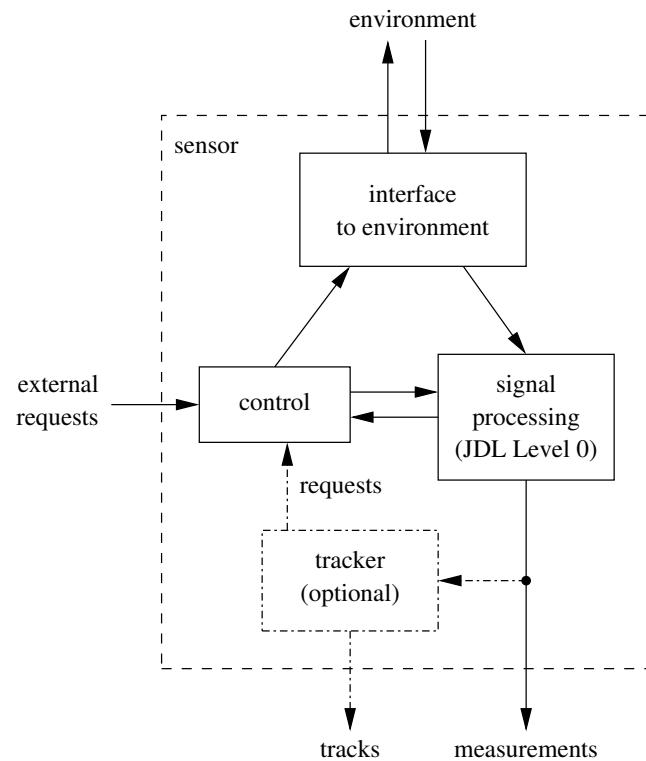


Figure 3: Sensor architecture.

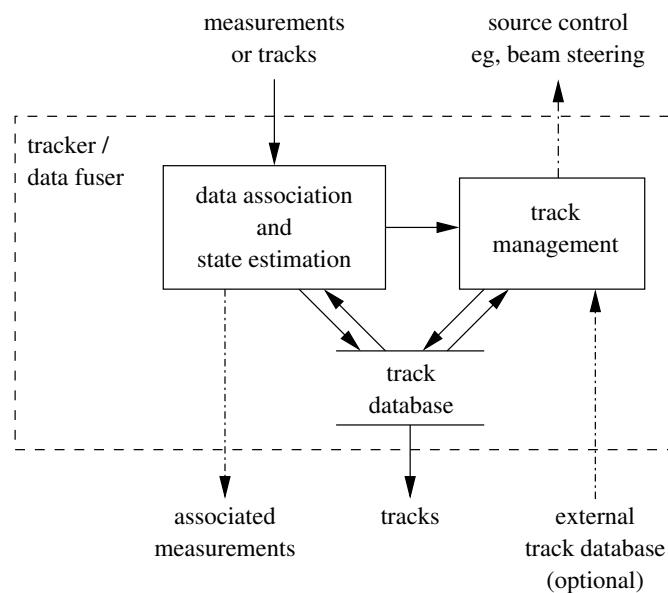


Figure 4: Tracker architecture.

In addition, subsystems may be specified and constructed by different manufacturers, or may contain legacy hardware with existing constraints. In general, the development of a distributed sensor fusion system needs to be a combination of ‘system engineering’ and ‘system architecting’ approaches [Waltz & Hall 2001]: broadly, the former supports a first principles design, with requirements that flow down to system specifications, and the latter seeks to optimise the usage of standard components. More details regarding the system engineering approach may be found in Bowman & Steinberg [2001].

A number of issues can be identified from this multiple platform, multiple sensor, distributed sensor fusion model. These are described in the following section.

3 Distributed Multisensor Fusion Issues

The following is a categorisation and description of multisensor fusion issues that arise when designing or upgrading a modern combat system. It is based on material from the open literature and experience with a range of Australian Defence acquisition projects. It expands on the data fusion node paradigm of Bowman & Steinberg [2001, Figure 16.5], which incorporates three basic functions: data alignment, data association and entity state estimation. A summary of the issues is provided as Appendix A.

1. Sensor data

In this document, sensor measurements² are considered to be the most fundamental type of information available for sensor fusion. The optimal data fusion approach requires measurements to be available to all platforms wishing to make use of the sensor’s data. Whether or not measurements are broadcast is a function of the architecture.

(a) sensor ID

The sensor should be uniquely identified, and there should be some indication of the type of sensor that generated the measurement [Shapell 1997, Section 3.9.6.3]. Data association and fusion algorithms require knowledge of the type and properties of the data.

(b) measurement data elements

Typical sensor measurements are target range, azimuth, elevation and Doppler velocity for primary radars, azimuth, range and identity for secondary radars, azimuth and elevation for passive electro-optic sensors, and azimuth and emitter class for electronic surveillance sensors. However, depending on the system architecture, not all measured data may be made available to sensor fusion systems. In addition, physical and system constraints may introduce other effects, such as the truncation (for example, limits for target elevation) and quantisation of data.

Supplementary information may include the signal-to-noise ratio and false plot probability. False plots are measurements that do not correspond to a target.

²Raw sensor measurements are also known as detections, contacts, plots or strobes (for bearing-only measurements).

The probability of a false plot occurring may be required by the data association algorithm.

The means used by the sensor to obtain the measurements may or may not be important to a tracking system. For example, a chirp waveform may exhibit a range-Doppler ambiguity that cannot be resolved without access to other data.

(c) measurement uncertainty

Measurements are inherently uncertain, that is, they contain errors, for a number of reasons including environmental and system noise, the finite resolution of sensors, and the physical characteristics of the target. The uncertainty of each measurement parameter is required at each network node by the sensor and data fusion algorithms. Note that sensor fusion does not compensate for poor sensors [Hall & Garga 1999].

(d) target revisit rate

This is the interval between consecutive measurements corresponding to a particular target. The revisit rate may be fixed, as in the case of a rotating mechanically scanned radar or scanning electro-optical sensor, or variable, as for a phased array radar.

(e) sensor data rate

This is a measure of the volume of data that needs to be accommodated by the system that handles the sensor data. It is a function of the types of data output by the sensor, such as raw measurements or tracks, the false alarm rate, the number of targets, and the target reporting priorities.

(f) sensor location and orientation

The location and orientation of the sensor may be reported absolutely or with respect to the platform carrying the sensor. There may also be a combination of the two, such as when a stabilised ship based radar provides bearing relative to the platform and elevation relative to the horizon. The location and orientation details of the platform may therefore be important. The uncertainty in all of these values should also be available.

2. Sensor registration

Sensor registration is the alignment of sensor data so that the reported positions of each target from multiple sensors correspond to the same physical and temporal locations. Sensor registration is critical to the performance of a multisensor tracking system; the fusion of data from two unregistered sensors may give results that are worse than the data from either sensor acting alone [Moore & Blair 2000].

(a) sensor error sources

The sensor parameters that are subject to biases include range, azimuth, elevation and time. Range may also have a scaling error, and azimuth and elevation bias errors may be time-varying due to the sensor motion.

(b) errors in sensor platform position and orientation

The determination of sensors' locations is crucial to the registration process, since a sensor's spherical measurements are specified relative to its position and orientation. The positions of networked sensors may be specified relative to one

another, or with respect to a reference location, such as the centre of the earth. With the availability of the Global Positioning System (GPS), it is practical for a platform's position to be specified absolutely to within the resolution of its sensors. Heading, pitch and roll are generally estimated and compensated for by the platform's navigation system. Any related errors may be manifested as time-varying bearing errors.

(c) relative and absolute registration

Relative sensor registration uses the local platform as a reference, and corrections are applied to remote data. It is simpler than absolute registration, where all sensor biases are estimated, but it may not give the performance that is desired when combining data from accurate sensors [Moore & Blair 2000, p. 52]. However, relative registration may have advantages when data are used with local systems, such as organic weapons.

(d) coordinate transformation errors

Errors can arise from computational limitations, such as 'round-off' errors and algorithmic approximations, when data are translated between coordinate systems. A good example is the translation from earth centred Cartesian coordinates to latitude, longitude and altitude [Best 2003], which does not have a closed form solution and relies upon approximations to obtain an iterative solution. However, the greatest contributors to the problem of coordinate transformation errors are inaccuracies in the earth model and the combination of multiple, two dimensional sensor measurements, which is an under-determined problem [Moore & Blair 2000, p. 50].

(e) time synchronisation

The synchronisation of data from different sources is critical, and may be a problem if the data transmissions are not time stamped. If measurements are transmitted to other platforms, timing errors may produce inconsistencies in the surveillance pictures among these participants. Timing is discussed further under Sensor Networking (4).

(f) uncertainties in the registration error estimates

Sensors cannot be perfectly registered. The magnitude of the residual errors contribute to track accuracy and are important when determining the accuracy of a networked sensor system.

3. Tracks

A track represents a platform's knowledge of a target, derived from the information from the on-board sensors, remote sources and the operator. Each track should be allocated an unique number. Tracks should also include the time of the most recent update, the target's country of origin, category, type, class, name, threat designator, position, course, speed, and kinematic uncertainty, other uncertainty information regarding the derivation of this information (such as via sensor fusion), data dependencies, security classification (of the fields, message or record), and comments [Shapell 1997, Appendix I]. Some track specific issues follow.

(a) track management

Track management is the process of maintaining a consistent track database that reflects the estimated state of the targets in the region of interest. This is difficult in network centric applications due to timing issues [Moore & Blair 2000].

i. track initiation

Track initiation may be based on data from individual sensors, individual platforms, or networked sensor data. The algorithm, and its data sources, will affect the probability of an early initiation of a track on a target, and the probability of a false track.

Consistency of track initiation algorithms is not necessary for picture consistency across network participants, even if initiation is based on networked sensor data. This is because it does not matter which participant initiates a track, only that participants agree on plot-to-track and track-to-track association once a track has been initiated. This may give considerable flexibility to the processing of individual sensor data to allow for the customising of track initiation according to platform requirements, sensor performance and environmental conditions.

ii. track number management

Tracks should be unambiguously identifiable across all network participants. This greatly simplifies the association of remote track reports.

iii. track deletion

Inconsistencies in track deletion rules across the network, such as the interval following a track update that is required for a track to be considered ‘lost’, will result in differing track pictures.

iv. track merging

Track merging is the recognition that two targets have become unresolvable by the network’s sensors, and the subsequent dropping one of the corresponding tracks. As per track association and deletion, it is important for picture consistency that track merging rules are consistent across the network.

v. track divergence

This is a special case of track initiation, where a new target appears with initial conditions that correspond to an existing target. An example of this is the launch of a missile by a fighter aircraft or the splitting of a raid formation. Unlike for track initiation, the rules for the detection of diverging targets need to be consistent across the network to avoid picture inconsistencies. This is because it is important that all participants recognise which track is associated with a given plot.

vi. conflict resolution

Inconsistencies may occur between the track pictures on different platforms, resulting in duplicate tracks or swapped track numbers, for example. There needs to be a mechanism for recognising and resolving inconsistencies. In the crudest of systems, these functions may be performed by human operators communicating via a voice channel.

(b) **track coordinate system(s)**

The coordinate system used for tracking may be chosen so that target motion is linear in that coordinate system. If linear target motion results in non-linear effects in the adopted coordinate system, there may need to be mechanisms to compensate. The frame(s) of reference for the coordinate system (for example sensor, platform or centre of the earth) may depend upon the application and the data provided by the sensor registration process [Moore & Blair 2000, p. 27]. Tracks may be two dimensional or three dimensional, depending on the state variables of interest and the information available from the sensor. Fusing tracks with different coordinate systems may be complex, but may also resolve ambiguities in other dimensions. For example, if the tracks are from separated radars that only measure range and azimuth, combining them may provide an estimate of target height.

(c) **target state representations**

Not all state parameters may be represented by the track data structure that is broadcast around the sensor network. In addition, data will be quantised to some level of precision, and may be truncated (for example, limits for target altitude). These will all need to be taken into account by the user of the data.

(d) **uncertainty representations**

The track data should include an estimate of the uncertainty in the target state estimate, for example the probability density function, covariance, part of the covariance matrix, or a figure of merit. The type of uncertainty information may determine the utility of the track data to the recipient: having only a figure of merit may be sufficient to support a Reporting Responsibility networking model, but may preclude fusing the track data with local data, for example.

(e) **target motion models**

i. **appropriate**

The motion models used by the system to predict target positions and velocities should be appropriate for the targets being tracked. For example, the task of tracking a fighter aircraft is quite different from tracking a ballistic target. Fighter aircraft are highly manoeuvrable, and multiple motion models may be required for straight and level flight, lateral turns, linear acceleration, climbs, dives, etc. A ballistic target is more predictable and is likely to follow a predefined trajectory that may be modelled with a single model. There may also be known upper or lower limits that can be applied to the speed or acceleration model for a target.

ii. **consistent**

Motion models should be consistent between participants sharing low level data. Inconsistencies in track predictions will lead to misassociations between plots and tracks, resulting in significant picture inconsistencies.

(f) **attributes**

Track attributes are non-kinematic aspects of target characteristics. They usually refer to identification related parameters, for example category (air, surface, subsurface, etc.), type (FFG, DDG, etc.), class (frigate, destroyer, etc.), threat designation (hostile, friendly, neutral, etc.), name (for example, HMAS

Perth), nationality, and a classification confidence [Shapell 1997, Section 3.9.4]. Other attributes may include fuel status, weapon status and other information of interest.

(g) **plot–track association**

There are many algorithms for selecting the measurements that contribute to a particular track, for example, nearest neighbour, probabilistic data association and multiple hypothesis tracking [Moore & Blair 2000, p. 54]. The choice of algorithm may depend on the target density, clutter, target type, available computational resources, latency restrictions and required tracking performance. As with target motion models, the issues are appropriateness for the task and consistency between network participants where plot messages are shared.

(h) **plot–track fusion**

The updating of a track with a measurement is fundamental to tracking. There are many algorithms for performing this function, including linear and Extended Kalman filters [Bar-Shalom & Li 1995], Unscented Filters [Julier & Uhlmann 2001b] and particle filters [Arulampalam, Maskell, Gordon, & Clapp 2002]. Again, the issues of target type, computational complexity, latency and tracking performance will influence algorithm selection. Algorithms should be consistent between network participants where plot messages are shared.

(i) **track–track association**

Track association is the recognition that two tracks represent the same target. It is critical that the network participants employ consistent rules regarding the association of tracks. Significant differences in track pictures will result from inconsistent association decisions.

A variety of algorithms are available for track–track association, including nearest neighbour, global nearest neighbour, and multiple hypothesis. The type, and effectiveness, of the algorithm may be influenced by the available information and the confidence in that information. For example, a data link may not provide the full track covariance matrix, or the transmitting source may artificially increase or decrease the reported uncertainty in the data.

(j) **track–track fusion**

It is possible that the system will combine tracks from local or remote sources to obtain a single track for each target. There are many algorithms for this purpose, ranging from the selection of the ‘best’ track, through Bayesian combination [Bar-Shalom & Li 1995], to covariance intersection [Julier & Uhlmann 2001a]. Practical issues include asynchronous sensors and other data sources, dependencies between tracks, track registration and the availability of sufficient error statistics.

(k) **false tracks**

False tracks are tracks that do not correspond to a target, that is, they are formed entirely from sensor reports that are not produced by real targets. It is necessary to define ‘false’ as opposed to ‘unwanted’ tracks, such as tracks on birds or swarms of insects, etc. Of interest are the rate at which they are formed, the mechanisms for identifying them and their duration. The prevalence of false tracks will be strongly influenced by the sensor network architecture, for

example, whether tracks are initiated on measurements from individual sensors or from measurements received from all sensors.

(l) **track data sources**

In addition to local and remote sensors, combat system track data may be sourced from operator inputs or databases. This is particularly applicable to attribute data that may not be available from sensors, such as target identification and nationality [Shapell 1997, Section 3.9.7.3]. Other ‘prior’ information may contribute to the track picture, including kinematic limits for specific target types, airspace control details, surface and air tasking orders, flight plans, safe corridors and shipping lanes, and electronic surveillance databases for emitter and platform identification.

(m) **track pedigree**

Data fusion systems should maintain a record of the sources that contribute to a fused data product. This will help avoid the inappropriate reuse of data, and it will assist operators with correcting inconsistencies in the situational picture and identifying faulty sensors and remote sources. Date reuse is discussed further under issue 4.(e).

(n) **attributes of fusion products compared with local sensor data**

The system may exhibit characteristics resulting from the sensor network that may be compared with the performance of an isolated sensor. An analysis of these are useful when determining the qualitative or quantitative value of networking the sensors. Examples are gains in kinematic accuracy, such as cross-range accuracy improvement through combining data from two separated sensors that have good range resolution, improved track continuity, improved robustness to jamming, and increased false track rate and duration. A sound design will strengthen the advantages of sensor networking while attempting to negate the effects of the disadvantages.

4. Sensor networking

This topic considers issues related to the combining of data in distributed sensor networks.

(a) **architecture**

There are many types of network architecture, ranging from centralised, where all data are passed to a single, central tracking system, to fully distributed, where each participant builds its own picture from data received from all participants, with many hybrid combinations in between. Each type has its advantages and disadvantages, with Moore & Blair [2000] being a useful reference for comparing different sensor fusion architectures.

(b) **communicated data**

Different types of sensor data may be communicated between platforms, depending on the architecture and algorithms employed in the network.

i. **measurements**

The communication of raw measurements provides optimal results, but at a significant cost in terms of communication and computational requirements [Moore & Blair 2000, pp. 10–11]. An alternative is to broadcast only

those measurements that are associated with tracks, that is, Associated Measurement Reports (AMRs) [Moore & Blair 2000, pp. 11–13]. Provided the network nodes use the same association and fusion algorithms, the track maintenance and tracking accuracy will be the same as if all measurements were broadcast. The primary difference is with track initiation, since tracks may only be initiated by data being available to a single platform. However, this will only differ from the optimal case where sensors on different platforms are involved in track initiation on the same target simultaneously, a rare occurrence in practice.

ii. **tracks**

Depending upon the update rate, the transmission of tracks may provide a saving in bandwidth, but at the expense of accuracy [Moore & Blair 2000, pp. 13–19]. This approach is simple, and the accuracy may be adequate for surveillance applications. There are issues regarding the assignment of reporting responsibility among the participants.

iii. **tracklets and other data summaries**

Tracklets provide a summary of measurements [Moore & Blair 2000, pp. 16–17], and may be used to update a track where update rates and accuracy are not critical. The update rate may easily be adjusted dynamically. Tracklets give the same results as using the individual measurements in applications where the motion of the target is predictable, such as with a trajectory that is known to be linear or ballistic. Tracks may also be represented using summaries of the state vector and/or error covariance matrix [Wang, Evans, Challa, Mušicki & Legg 2003]. Although these schemes are suboptimal, since the original data cannot be fully recovered, they may allow reasonable data fusion performance using a fraction of the network bandwidth of optimal approaches.

iv. **fusion products**

Some data links only allow the transmission of data that have been derived from local sensors. It may be useful to broadcast the data that result from the fusion of local and remote sensors, however the sources of the fused information should be maintained [Shapell 1997, Section 3.9.7].

v. **data push/pull**

‘Data push’ refers to the transmission of data that has not been requested, and ‘data pull’ refers to the receiving of requested data [Shapell 1997, Section 3.9.9]. The former is common among data distribution systems, but may be an inefficient use of processing capability and communication channel bandwidth. The latter may require changes that are doctrinal, as well as technical, and may be supported by an internet-like protocol. To optimise the usage of a data channel with a finite bandwidth or latency, a flexible surveillance data distribution system may need to know the requirements of the users of the received data in order to distribute the necessary data at the appropriate accuracy. For example, a platform may have an excellent 2-D accuracy and require height from a third party.

(c) sensor data releasability

The networking of sensor data necessitates the exposure of the low level performance of the participating sensors to all users of the data. This may be an issue where sensor manufacturers do not wish to reveal actual equipment performance to other nations or manufacturers. Although this performance may be disguised through the use of tracklets, say, rather than raw measurements, recipients of the data still require some assurance that the report is valid.

(d) track reporting responsibility

The Reporting Responsibility (R2) approach to maintaining a common surveillance picture requires each track to be broadcast by only one platform, preferably the one that has the most accurate estimate of the target's state [Moore & Blair 2000, pp. 7–10]. Each platform maintains local tracks, and it assumes R2 when it determines that its accuracy is superior to that of the broadcasting platform. This is facilitated using a figure of merit known as Track Quality (TQ) [Shapell 1997, Sections 3.9.5.3–3.9.5.6].

Track Quality and the rules of Reporting Responsibility achieve four objectives [Shapell 1997, Section 3.9.5.3]:

- i. they ensure that each track is reported by only one participant,
- ii. they establish which participant reports each track,
- iii. they give R2 to the platform that achieves the greatest accuracy, and
- iv. Track Quality indicates the accuracy of the track.

(e) input from other data links

In addition to data derived from participants' sensors, sensor networks can incorporate data from other data links. The management of two-way surveillance data exchanges may involve complex data translation, track management and, possibly, data reuse issues.

The translation of data from one format to another may be nontrivial owing to the differing requirements and philosophies behind the data links. For example, it may be complex (and even undesirable) to translate data between a high bandwidth, high accuracy network and a low bandwidth, wide area data distribution system.

The inadvertent reuse of data, or data incest, results in tracks that have an error that is increased, but a reported uncertainty that is erroneously decreased [McLaughlin, Evans & Krishnamurthy 2003]. This condition may be avoided by the careful management of data sources, or by compensating for the data reuse. In practice, however, the effect may be small.

(f) capacity issues

The required capacity of a sensor network is influenced by a variety of issues, such as the number of targets to be tracked, the required tracking accuracy, the sensor data representation, the data reporting protocol and the number of network participants.

(g) available bandwidth

The volume of data that may be distributed across a network may be limited by the data distribution hardware, and has a strong influence on the architecture

of the data distribution system and the delays between track updates. It may be necessary for participants to prioritise the data to be broadcast.

(h) **measurement/track data update rate**

The track estimate update rate may be dependent upon the data rate of the contributing sensor(s), the number of targets being tracked and their relative priorities, and the available communications bandwidth. The data rate will impact the track accuracy and continuity.

(i) **latency**

Delays between the sensor detections, or tracks being updated, and the data being received by other participants may cause problems with the consistency of measurement or track data across a network, and, subsequently, with track management. The delays may be deterministic (such as those that are dependent upon the design of the network) or stochastic (for example, being dependent upon the real time track priorities). Multi-source data are frequently asynchronous, and time stamping is essential.

(j) **out-of-sequence data**

It may be possible for associated measurements or track updates to be received following the arrival of more recent data. A strategy may need to be in place to deal with this, such as to combine it optimally [Challa, Evans & Wang 2003], or discard it if the data are older than a predetermined threshold [Moore & Blair 2000, pp. 43–44].

(k) **priority**

A track prioritising mechanism may be necessary when the bandwidth limit is approached [Shapell 1997, Section 3.9.7.5]. The priority of a track may be affected by the proximity of the track to a sensor, the status or identity of the track, and the use to which the track is put.

(l) **consistency between participants' pictures**

Users of surveillance data may have differing accuracy requirements that may (or may not) be met by the data distribution system. However, the users' surveillance pictures should be consistent, that is, there should be a direct correspondence between the tracks in each picture, including track numbers and identification.

(m) **network infrastructure**

This item is concerned with issues relating to the infrastructure of the sensor network.

i. **network scalability / the addition and removal of participants**

It is advantageous for networks to be scalable, that is, capable of growing or shrinking in the number of participants. In particular, participants should be able to join and leave the network without unnecessarily impacting network performance. Special consideration may be necessary when participants performing controlling functions leave the network.

ii. **robustness**

This is the extent to which the network can adapt if communications are lost between participants; it includes survivability and fault tolerance [Shapell 1997, Section 3.8].

iii. surveillance picture releasability

The network data distribution system may require a mechanism to prevent some participants from receiving all of the data for reasons of national security. The system could filter out sensitive tracks or track attributes and/or provide kinematic data at a lower accuracy than that at which it is capable.

iv. security

This refers to the network's resistance to detection or eavesdropping. It may be influenced by message encoding or the use of low power, spread spectrum transmissions with a low probability of intercept.

v. vulnerability

Vulnerability refers to the network's resistance to adverse environmental effects and jamming.

vi. common operating environment

Having common software across platforms for processing the data is beneficial from the perspective of software development and maintainability. A user interface that varies between participants complicates training.

(n) external communications

The surveillance picture may be a contributor to other, possibly wider area, surveillance picture distribution networks, or it may receive data from an external network. This may introduce hardware compatibility or track management issues that need to be addressed, for example, differences between track association algorithms that may result in duplicate tracks.

Data releasability will need to be considered where there is the potential for sensitive data, such as sensor performance, being compromised.

5. Sensor control

It may be possible for sensors to be controlled via the surveillance network. Such control may include the following. Basic emission control functionality, such as allowing a sensor to radiate, is exerted by the command and control system and is not discussed here.

(a) cuing

A sensor is directed to increase its probability of detection in a defined region so that it may detect a known target. This may allow a platform to obtain an accurate fix, or initiate a track, on a target sooner than it would otherwise be able to. Examples include cuing radars from electronic surveillance detections and cuing fire control radars from surveillance radars.

(b) detection threshold

As a result of the knowledge available to other sensors, the threshold used by a sensor to detect targets in noise may be changed to improve the detection probability in the vicinity of a known target.

(c) resource allocation

A sensor such as a phased array radar may have control over the allocation of its time and energy budget, so it can adaptively control its resources across

tasks such as horizon and long range air search, track updates and weapons control. Moore & Blair [2000] describe a scheme whereby a tracking filter can determine the time at which a new measurement is required from a phased array radar based on tracking parameters, such as the track accuracy, target manoeuvre and missed detections. This may result in a 50% saving in sensor resources over a conventional system.

(d) local/global control

Sensor control or management may be optimised locally at the sensor subsystem level, or globally at the track picture level. Phased array radars may use a hybrid of both, utilising the low latency (typically <10 ms) local control for detection revisits during track initiation, and the slower (typically >1 s) control for track revisits.

6. Hardware and software

This item is concerned with the physical aspects of the system and its implementation. This topic is discussed further by Moore & Blair [2000, p. 68].

(a) computing capability

The required capability of the computing hardware may be considered with respect to the system requirements, which may be considerable when there are complex tracking algorithms, a large number of targets, or when timeliness is critical, as with weapons control. It may be preferable for systems to have a capacity significantly in excess of that required to allow for future growth. The physical constraints on size, weight, power and cooling requirements may limit the available computing resources.

(b) initialisation time

The communications system may require a significant period of time to become operational when it is initialised. This may be an issue in the case of processor handover resulting from computer failure.

(c) personnel issues

i. personnel required to operate or maintain the system

The number and skills of personnel associated with the system may be a consideration, and may have a significant influence on the design of the system.

ii. human-machine interface

An objective of the system that distributes surveillance data is situation awareness, a state of mind of the end user of the system. This is facilitated (or hampered) by the human-machine interface of the picture dissemination system. A useful list of requirements for command and control system operator controls is given in [Shapel 1997, Annex 2A]. The training of system users is important, but may not be an issue for a technical review.

iii. maintenance and reliability

Areas of concern here include the mean time between failures and the mean time to repair. The ease with which hardware and software can be added or modified is also a consideration in the design.

iv. contractor issues

The ability of a contractor to deliver to the contracted performance, schedule and cost is a significant risk for any complex acquisition, particularly for tracking and sensor fusion systems where real sensor data may not be available until late in the acquisition schedule. Other significant issues include the availability and distribution of necessary and accurate information to sub-contractors, and the integration of sub-contracted components.

7. Performance specification and assessment

For the purposes of specifying requirements and assessing against these requirements, it is necessary to quantify the capability of the networked sensor system. Careful consideration of the purpose of the system is necessary to identify quantitative measures that are relevant. For example, in an anti-ship missile defence scenario, track initiation delay and track continuity may be more important than kinematic accuracy. It is also important to specify the conditions under which the performance measures apply.

Broadly, performance measures are classified as Measures of Performance (MOPs), which quantify a characteristic of a sensor fusion system, such as the false alarm rate, Measures of Effectiveness, which quantify the utility of the system with respect to operational considerations, such as the warning time available to a platform that is under attack, and Measures of Force Effectiveness, which quantify the ability to complete a mission [Llinas 2001]. Owing to the variations in system objectives, it is impossible to list all possible metrics; some useful references are Llinas [2001], who also discusses test and evaluation processes, Li & Zhao [2001], who discuss MOPs for estimators, and Colegrove, Davis & Davey [1996], who describe a tool for assessing tracking systems.

One complication with the quantitative assessment of sensor system tracks is access to an independent and accurate record of the targets' true dynamics, or ground truth. In practice, such a record may be erroneous or unavailable. The latter situation is discussed by Blackman & Dempster [2002], who consider performance metrics where targets of opportunity are used to assess tracking performance. They define five track categories, ranging from 'long and clean' to 'junk', and specify ad hoc metrics that categorise tracks. The metrics listed here assume access to independent data that provides the 'assumed truth' for each track [Colegrove et al. 1996].

In general, target states and sensor outputs are modelled as stochastic processes. Therefore, most tracking and sensor fusion related metrics are statistical in nature, and they are usually defined in terms of probabilities or statistical parameters (for example, the mean).

(a) track initiation

Track initiation refers to the formation of a new track. Typical track initiation metrics include the probability of track initiation, the track initiation delay (from the first detection) and the track initiation range.

(b) track maintenance

Track maintenance refers to the ability of the tracker to continue to estimate the parameters of the target. Example track maintenance metrics are the prob-

ability of track maintenance, track overshoot, the number of divergent tracks (that is, tracks that have an inappropriately low uncertainty given their distance from the target), the number of track swaps (where the track-target assignment changes), and the number of different tracks associated with a target track.

(c) **kinematic accuracy**

Kinematic accuracy metrics include absolute position, speed and heading errors, and azimuth, range and elevation errors from a specified position, such as a sensor. Since the tracker typically maintains a state estimate and covariance representing a Gaussian probability density function, these metrics need to be determined from data that have been translated into the appropriate coordinate system. In general, finding the range error based on the state in Cartesian coordinates will give a biased result; however, unbiased coordinate conversions [Longbin, Xiaoquan, Yiyu, Kang & Bar-Shalom 1998] or other techniques, such as the unscented transform [Julier & Uhlmann 2001b]), may be used.

Accuracy values should be specified in probabilistic terms, for example azimuthal error smaller than a specified value 90% of the time. Percentiles are useful for indicating the spread of a one-sided absolute error, and are far more meaningful than standard deviations since the distribution will not be Gaussian. In this case, outliers will exaggerate the standard deviation, but have little effect on the percentiles.

(d) **false tracks**

False track metrics include the frequency of false tracks and their average duration. False track performance is generally achieved at the expense of track initiation performance. It is important to discriminate between false tracks (based on false detections that arise from noise and environmental effects) and ‘unwanted’ tracks (which are tracks on real objects that are not of interest, such as birds and insect swarms).

(e) **identity veracity**

It is important that track identity estimates are accurate, be they category, type, class or threat designation [Shapell 1997, Section 2.8.8]. This is critical to situation awareness. The time taken for an identification may also be important.

(f) **timeliness**

The interval from the time of a sensor detection to the time a track or other entity is updated on a display, passed for further processing or transmitted is crucial to effective operational performance. The allowed latency will be influenced by acceptable delays for operators and time constraints for other functions, such as fire control.

(g) **picture completeness**

Appropriate metrics for representing the completeness of the surveillance picture are the number of omitted tracks, the number of false tracks and the number of duplicated tracks. These are similar to some of the track maintenance metrics.

(h) **consistency between participants' pictures**

The consistency between the pictures available to different participants should be confirmed. The acceptable differences between participants' pictures will depend on the application.

There are mechanisms for combining performance measures into overall figures of merit, for example, forming a weighted sum after assigning relative weights that reflect the priorities of the application [Colegrove et al. 1996].

4 Conclusions

This report has presented a collection of factors that may be considered when the tracking and sensor fusion aspects of a distributed surveillance system are specified or evaluated. These factors cover sensor data processing and distribution, tracking, networking issues, sensor control, computing resources, and performance specification and assessment. The relative priorities of each of these will depend upon the specific application and the implemented solution.

5 Acknowledgements

The author thanks Dr Mark Krieg for his helpful suggestions.

References

Arulampalam, M. S., Maskell, S., Gordon, N. & Clapp, T. (2002) A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking, *IEEE Transactions on Signal Processing* **50**(2), pp. 174–88.

Shapell, J., ed. (1997) *Handbook 5: Guidelines for Maritime Information Management*, AUS-CAN-NZ-UK-US Command, Control and Communications Committee, Iss. 2.1.

Bar-Shalom, Y. & Li, X.-R. (1995) *Multitarget-Multisensor Tracking: Principles and Techniques*, 3rd edn, Yaakov Bar-Shalom (YBS), Storrs, CT. ISBN 0-9648312-0-1.

Best, J. P. (2003) *Coordinate Usage in the Virtual Maritime System Architecture*, Virtual Maritime System document 00046, DSTO.

Blackman, S. & Dempster, R. (2002) Real-world tracking system performance evaluation, in O. E. Drummond, ed., *Signal and Data Processing of Small Targets: Proceedings of SPIE*, Vol. 4728, pp. 470–6.

Bowman, C. L. & Steinberg, A. N. (2001) A systems engineering approach for implementing data fusion systems, in Hall & Llinas 2001, Chapter 16.

Challa, S., Evans, R. & Wang, X. (2003) A Bayesian solution to the OOSM problem and its approximations, *Journal of Information Fusion* **4**(3).

Colegrave, S. B., Davis, L. M. & Davey, S. J. (1996) Performance assessment of tracking systems, in *Proceedings of the International Symposium on Signal Processing and its Applications (ISSPA)*, Signal Processing Research Centre, QUT, Brisbane, Australia, pp. 188–91.

Hall, D. L. & Garga, A. K. (1999) Pitfalls in data fusion (and how to avoid them), in *Proceedings of Fusion 1999*, ISIF, pp. 429–36.

Hall, D. L. & Llinas, J., eds (2001) *Handbook of Multisensor Data Fusion*, CRC Press, Boca Raton, FL.

Julier, S. & Uhlmann, J. K. (2001) Generalised decentralised data fusion with covariance intersection, in Hall & Llinas 2001, Chapter 12.

Julier, S. & Uhlmann, J. K. (2001) Data fusion in nonlinear systems, in Hall & Llinas 2001, Chapter 13.

Li, X. R. & Zhao, Z. (2001) Measures of performance for evaluation of estimators and filters, in *Signal and Data Processing of Small Targets: Proceedings of SPIE*, Vol. 4473, pp. 530–41.

Llinas, J. (2001) Assessing the performance of multisensor fusion processes, in Hall & Llinas 2001, Chapter 20.

Longbin, M., Xiaoquan, S., Yiyu, Z., Kang, S. Z. & Bar-Shalom, Y. (1998) Unbiased converted measurements for tracking, *IEEE Transactions on Aerospace and Electronic Systems* **34**(3), pp. 1023–7.

McLaughlin, S. P., Evans, R. J. & Krishnamurthy, V. (2003) Data incest removal in a survivable estimation fusion architecture, in *Proceedings of Fusion 2003*, ISIF, pp. 229–36.

Moore, J. & Blair, W. (2000) Practical aspects of multisensor tracking, in Y. Bar-Shalom & W. D. Blair, eds, *Multitarget-Multisensor Tracking Applications and Advances*, Vol. III, Artech House, Boston, MA, Chapter 1.

Steinberg, A. N. & Bowman, C. L. (2001) Revisions to the JDL data fusion model, in Hall & Llinas 2001, Chapter 2.

Waltz, E. & Hall, D. L. (2001) Requirements derivation for data fusion systems, in Hall & Llinas 2001, Chapter 15.

Wang, X., Evans, R., Challa, S., Mušicki, D. & Legg, J. (2003) *Distributed Sensor Fusion with Network Constraints*, Technical report, Cooperative Research Centre for Sensor Signal and Information Processing. Report to DSTO through the Tracking and Data Fusion Laboratory.

Appendix A Summary of Issues

Table A1: Summary of distributed multisensor fusion issues.

<p>1. Sensor data Issues related to the individual sensors, such as the type of data provided and its accuracy. (Sensor control is addressed separately.)</p>	<ul style="list-style-type: none"> (a) sensor ID (b) measurement data elements (c) measurement uncertainty (d) target revisit rate (e) sensor data rate (f) sensor location and orientation
<p>2. Sensor registration Sensor registration, or ‘grid-locking’, determines the relationship between different sensor coordinate systems, so that data may be usefully combined.</p>	<ul style="list-style-type: none"> (a) sensor error sources (b) errors in sensor platform position and orientation (c) relative and absolute registration (d) coordinate transformation errors (e) time synchronisation (f) uncertainties in the registration error estimates
<p>3. Tracks Issues related to processed local and/or remote data, generally represented as tracks.</p>	<ul style="list-style-type: none"> (a) track management <ul style="list-style-type: none"> i. track initiation ii. track number management iii. track deletion iv. track merging v. track divergence vi. conflict resolution (b) track coordinate system(s) (c) target state representations (d) uncertainty representations (e) target motion models <ul style="list-style-type: none"> i. appropriate ii. consistent (f) attributes (g) plot-track association (h) plot-track fusion (i) track-track association (j) track-track fusion (k) false tracks (l) track data sources (m) track pedigree (n) attributes of fusion products compared with local sensor data
<p>4. Sensor networking Issues related to the communication aspects of sharing sensor data, such as the required bandwidth.</p>	<ul style="list-style-type: none"> (a) cuing (b) detection threshold (c) resource allocation (d) local/global control

Table A1: Summary of distributed multisensor fusion issues (cont.).

<p>5. Sensor control It may be desirable for the behaviour or performance of sensors to be under the control of the tracking system.</p>	<ul style="list-style-type: none"> (a) architecture (b) communicated data <ul style="list-style-type: none"> i. measurements ii. tracks iii. tracklets and other data summaries iv. fusion products v. data push/pull (c) sensor data releasability (d) track reporting responsibility (e) input from other data links (f) capacity issues (g) available bandwidth (h) measurement/track data update rate (i) latency (j) out-of-sequence data (k) priority (l) consistency between participants' pictures (m) network infrastructure <ul style="list-style-type: none"> i. network scalability / the addition and removal of participants ii. robustness iii. surveillance picture releasability iv. security v. vulnerability vi. common operating environment (n) external communications
<p>6. Hardware and software The physical aspects of the system.</p>	<ul style="list-style-type: none"> (a) computing capability (b) initialisation time (c) personnel issues <ul style="list-style-type: none"> i. personnel required to operate or maintain the system ii. human-machine interface (d) maintenance and reliability (e) contractor issues
<p>7. Performance specification and assessment Quantifying the relevant performance aspects of a networked sensor system is important and nontrivial.</p>	<ul style="list-style-type: none"> (a) track initiation (b) track maintenance (c) kinematic accuracy (d) false tracks (e) identity veracity (f) timeliness (g) picture completeness (h) consistency between participants' pictures

DISTRIBUTION LIST

Distributed Multisensor Fusion System Specification and Evaluation Issues

Jonathan A. Legg

	Number of Copies
DEFENCE ORGANISATION	
Task Sponsor	
Director General Maritime Development	1 (printed)
S&T Program	
Chief Defence Scientist	1
FAS Science Policy	
AS Science Corporate Management	
Director General Science Policy Development	
Counsellor, Defence Science, London	
Counsellor, Defence Science, Washington	
Scientific Adviser to MRDC, Thailand	
Scientific Adviser Joint	
Navy Scientific Adviser	
Scientific Adviser, Army	Doc Data Sheet
Scientific Adviser to the DMO	Doc Data Sheet
Information Sciences Laboratory	
Chief, ISR Division	Doc Data Sheet and Dist List
Research Leader, Wide Area Surveillance	Doc Data Sheet and Dist List
Head, Tracking and Sensor Fusion	1
Jonathan Legg	1 (printed)
Mark Krieg	1
Branko Ristic	1
DSTO Library and Archives	
Library, Edinburgh	1 and Doc Data Sheet
Defence Archives	1
Capability Development Group	
Director General Capability and Plans	Doc Data Sheet
Assistant Secretary Investment Analysis	Doc Data Sheet
Director Capability Plans and Programming	Doc Data Sheet

Director Australian Defence Simulation Office	Doc Data Sheet
Chief Information Officer Group	
Deputy Chief Information Officer	Doc Data Sheet
Director General Information Policy and Plans	Doc Data Sheet
AS Information Strategy and Futures	Doc Data Sheet
AS Information Architecture and Management	Doc Data Sheet
Director General Information Services	Doc Data Sheet
Strategy Group	
Director General Military Strategy	Doc Data Sheet
Assistant Secretary Strategic Policy	Doc Data Sheet
Assistant Secretary Governance and Counter-Proliferation	Doc Data Sheet
Navy	
Deputy Director (Operations),	
Deputy Director (Analysis),	
Maritime Operational Analysis Centre, Building 89/90 Garden Island Sydney NSW	Doc Data Sheet and Dist List
Director General Navy Capability, Performance and Plans, Navy Headquarters	Doc Data Sheet
Director General Navy Strategic Policy and Futures, Navy Headquarters	Doc Data Sheet
Air Force	
SO (Science), Headquarters Air Combat Group, RAAF Base, Williamtown	Doc Data Sheet and Exec Summ
Army	
ABCA National Standardisation Officer, Land Warfare Development Sector, Puckapunyal	Doc Data Sheet
SO (Science), Deployable Joint Force Headquarters (DJFHQ)(L), Enoggera QLD	Doc Data Sheet
SO (Science), Land Headquarters (LHQ), Victoria Barracks, NSW	Doc Data Sheet and Exec Summ
Joint Operations Command	
Director General Joint Operations	Doc Data Sheet
Chief of Staff Headquarters Joint Operation Command	Doc Data Sheet
Commandant ADF Warfare Centre	Doc Data Sheet
Director General Strategic Logistics	Doc Data Sheet
COS Australian Defence College	Doc Data Sheet
Intelligence and Security Group	
AS Concepts, Capability and Resources	Doc Data Sheet

DGSTA, Defence Intelligence Organisation	1 (printed)
Manager, Information Centre, Defence Intelligence Organisation	1
Assistant Secretary Capability Provisioning	Doc Data Sheet
Assistant Secretary Capability and Systems	Doc Data Sheet
Defence Materiel Organisation	
Deputy CEO, DMO	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet
Head Electronic and Weapon Systems Division	Doc Data Sheet
Program Manager Air Warfare Destroyer	Doc Data Sheet
Defence Libraries	
Library Manager, DLS-Canberra	Doc Data Sheet
OTHER ORGANISATIONS	
National Library of Australia	1
NASA (Canberra)	1
UNIVERSITIES AND COLLEGES	
Australian Defence Force Academy Library	1
Head of Aerospace and Mechanical Engineering, ADFA	1
Hargrave Library, Monash University	Doc Data Sheet
INTERNATIONAL DEFENCE INFORMATION CENTRES	
US - Defense Technical Information Center	1
UK - Dstl Knowledge Services	1
Canada - Defence Research Directorate R&D Knowledge and Information Management (DRDKIM)	1
NZ - Defence Information Centre	1
ABSTRACTING AND INFORMATION ORGANISATIONS	
Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Scientific Abstracts, US	1
Documents Librarian, The Center for Research Libraries, US	1
SPARES	
DSTO Edinburgh Library	5 (printed)

Total number of copies: printed 8, pdf 21

Page classification: UNCLASSIFIED

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. CAVEAT/PRIVACY MARKING		
2. TITLE Distributed Multisensor Fusion System Specification and Evaluation Issues		3. SECURITY CLASSIFICATION Document (U) Title (U) Abstract (U)		
4. AUTHORS Jonathan A. Legg		5. CORPORATE AUTHOR Defence Science and Technology Organisation PO Box 1500 Edinburgh, South Australia 5111, Australia		
6a. DSTO NUMBER DSTO-TN-0663	6b. AR NUMBER 012-918	6c. TYPE OF REPORT Technical Note	7. DOCUMENT DATE October, 2005	
8. FILE NUMBER E9505/25/115	9. TASK NUMBER NAV 02/152	10. SPONSOR DGMD	11. No OF PAGES 21	12. No OF REFS 20
13. URL OF ELECTRONIC VERSION http://www.dsto.defence.gov.au/corporate/reports/DSTO-TN-0663.pdf		14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved For Public Release</i> OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SOUTH AUSTRALIA 5111				
16. DELIBERATE ANNOUNCEMENT No Limitations				
17. CITATION IN OTHER DOCUMENTS No Limitations				
18. DSTO RESEARCH LIBRARY THESAURUS multisensors, sensor fusion, evaluation, specifications				
19. ABSTRACT This report discusses the tracking and sensor fusion issues that should be addressed during the specification or evaluation of a distributed, multisensor surveillance system for the provision of a common situational awareness picture or remote weapons control data. These issues include low level sensor data, sensor registration, target state estimation, sensor networking, sensor control, hardware and software, and performance specification and assessment.				

Page classification: UNCLASSIFIED

